Dr. James Peery, Director of the Information Systems and Analysis Center
Sandia National Laboratories

Committee: Senate Armed Services Committee,
Subcommittee on Emerging Threats and Capabilities

Subject: "Emerging Threats in Cybersecurity R&D"

Testimony:

Statement of Dr. James Peery, Director of the Information Systems and Analysis Center, Sandia National Laboratories

Senate Armed Services Committee, Subcommittee on Emerging Threats and Capabilities

March 20, 2012

Introduction

Chairman Hagan, Ranking Member Portman, and distinguished members of the Senate Armed Services Committee, thank you for the opportunity to testify. I am James Peery, Director of the Information Systems and Analysis Center at Sandia National Laboratories. Sandia is a multi-program national security laboratory owned by the United States Government and operated by Sandia Corporation for the National Nuclear Security Administration (NNSA).

Sandia is one of the three NNSA laboratories with responsibility for stockpile stewardship and annual assessment of the nation`s nuclear weapons. Within the U.S. nuclear weapons complex, Sandia is uniquely responsible for the systems engineering and integration of nuclear weapons in the stockpile and for the design, development, and qualification of all non-nuclear components of nuclear weapons. While nuclear weapons remain Sandia`s core mission, the science, technology, and engineering capabilities required to support this mission position us to support other aspects of national security as well. Indeed, there is natural, increasingly significant synergy between our core mission and our broader national security work. This broader role involves research and development in nonproliferation, counter proliferation, counterterrorism, energy security, defense, and homeland security. With the United States growing dependence on information technology, cyber security has become a key foundation in all of these areas.

Sandia's extensive cyber research and development (R&D) program is rooted in its rich history of providing adversarial threat assessments for the U.S. nuclear command and control systems. This program draws heavily upon our core science and technology (S&T) capabilities. These S&T investments afford the nation the ability to leverage world-leading capabilities in advanced analytics, trusted microelectronics, and modeling and simulation. Sandia's differentiating value comes from its unique systems approach

integrating scientific understanding, technology development, and complex requirements-driven engineering to develop solutions.

Sandia has developed a comprehensive understanding of mission needs and constraints through its long-standing relationship with key government agencies. Working in partnership with government, other national laboratories, academia, and industry, Sandia has been a key to:

- Providing technical leadership in threat-informed information assurance technology development and assessment
- Serving as an operational model for information security – with a goal of defining effective operational security guidelines and practice for Sandia, other government agencies, and high-value private-sector networks
- Expanding the cadre of highly skilled cyber professionals through its hands-on research internship program
- Functioning as a hub that works at the intersection of academia, national laboratories, industry, and government to drive cyber innovation and advance the overall national and global cyber health

My statement today will focus on a number of the challenges and technical developments in cyber security along with how the Department of Energy (DOE) laboratories contribute to the Department of Defense (DoD) mission in cyber security. I have been employed within the DOE labs for 22 years collectively, 17 of those years at Sandia National Laboratories, where I have done research in high performance computing and high energy density physics. Within management, I have led teams in cyber security, computational physics, high performance computing, nuclear weapons R&D and hydrodynamic testing. For the past two years, it has been my privilege to lead the organization at Sandia that represents the largest collection of cyber experts within the DOE laboratories. My testimony represents the vast knowledge that they have imparted to me.

Major points of this testimony

It is the belief of a Sandia team of cyber security experts that:

1. The DOE laboratories are a resource to DoD in "raising the bar" to the adversaries in cyber security. **We believe that a large part of the DoD is aware of where the cyber talent resides within the DOE laboratories and has effectively used DOE procedures to acquire that talent.**

2. A silver bullet for solving the "cyber problem" for DoD, DOE, dot-gov or the private sector does not exist. It is impossible to make an absolutely secure information technology (IT) system. **Sustained and coordinated investment in and deployment of government-owned science and technology could dramatically change the cost equation for our adversaries.**

3. **Compliance-based security and attempting to secure the perimeter are not effective.** We need a set of metrics to objectively measure system security. New technologies and policies should be evaluated and adopted based on how they objectively improve system security and how much they cost. This is not a static process as adversaries also adapt.

Based on the Committee's request, the following topics are addressed:

1. Mechanisms to rapidly develop, test, and field innovative approaches to address the expanding threat spectrum
2. Research on network security versus data encryption
3. Research on the transition from signature-based detection of attacks to behavioral detection
4. Test and evaluation infrastructures at various classification levels (e.g. digital sandboxes)
5. Other research priorities
6. Workforce issues
7. Coordination across the community

More can be said about these topics in a closed session.

1. Mechanisms to rapidly develop, test, and field innovative approaches to address the expanding threat spectrum: This issue is particularly relevant in the cyber domain, given the rate of change of both technology and threats. Historically, national security technology has evolved on the time scales of years. In the cyber realm, new exploits can render defenses that seemed effective obsolete in a matter of seconds. Given the speed with which cyber capabilities can be created and the relatively low cost for entry, the potential for possibly far-reaching technological surprise is very high.

Technology innovation has two key components: creation and adoption. One can support technology creation by providing consistent funding to create and maintain effective facilities and to attract properly trained researchers who are immersed in the problems of the day. Positive and open competition can be a powerful incentive to operate efficiently. I spent more than a decade of my career in the NNSA Advanced Simulation and Computing (ASC) program. Its goals were clear and technically compelling, we had challenging milestones, and funding was relatively stable. Because of those government investments, today we certify the U.S. nuclear weapon stockpile without the need for underground testing. Overall, the ASC program should be considered both an enormous technical success and a government success for a critical national security problem.

Creating a new technology and getting it adopted are two different tasks. There are significant barriers that prevent technology adoption including expediency, cultural inertia, and investments in legacy technologies. The business case for investing in new security technologies is often not clear, reinforcing the need for better metrics, risk assessment, and cost analysis.

Technology adoption can be accelerated by ensuring that researchers are partnered with users who understand operational needs and with vendors who can rapidly commercialize promising technology. Integrating and funding operational pilots as part of R&D programs can also improve the likelihood and pace of adoption. Results obtained from lab experiments are typically not enough to convince operators to deploy new technology. They need to see results in real world environments.

2. Research on network security versus data encryption: Encryption and network securities are complementary topics and should not be viewed as competing alternatives. Data encryption raises the bar for an adversary, but it is wrong to believe that encrypting all network traffic and all data at rest is sufficient to provide adequate security if you cannot also keep an enemy out of your networks. Again, there is no silver bullet. Our goal should be to raise the cost of successful attacks. Better network security and careful use of high quality encryption both raise adversary costs.

Cryptography is a based on well-understood mathematics. Time-tested algorithms and protocols exist. We can estimate how much work is required to break a given encryption scheme. Techniques exist for analyzing the security of cryptographic protocols. However, cryptography is quite subtle and it is easy to make mistakes especially in implementation. The early implementers of wireless communication protocols, who were all skilled engineers made numerous cryptographic errors. As technology evolves, effort is required to adapt the large body of cryptographic knowledge to the new technology. The adaptation is often straightforward and more of an engineering exercise than a basic research task.

Other aspects of network security are much less mature. For example, network filtering is often driven more by existing network protocols and recent exploitations than a coherent protection philosophy. Most networks use Transmission Control Protocol/Internet Protocol (TCP/IP) and thus base protection on filtering of TCP/IP packets, so filtering is limited to attributes visible in TCP/IP. Since TCP/IP has no notion of user identity, even a simple policy like "only administrators can configure the domain controller" requires multiple security mechanisms. A network filtering policy may ensure that only certain ports are open and that only certain types of packets can be sent to those ports. A host-based policy then ensures that only administrators have access to powerful configuration features. Verifying that this collection of policies properly enforces the desired abstract policy is difficult.

3. Research on the transition from signature-based detection of attacks to behavioral detection: Computer attacks have historically been detected using either signature- or anomaly-based methods. Anomaly-based techniques look for statistically significant deviations from normal activity. Because of the challenges in characterizing an accurate baseline of normal activity, anomaly-based detection systems to date have had limited utility. Signature-based methods, in contrast, compare network and file data against a database of known attack signatures to detect attempted intrusions and malware. Signature-based methods are incapable of detecting new attacks. Polymorphic malware

that can change its structure while retaining the same functionality is mostly immune to signature-based techniques.

More recently, a new class of anomaly detection methods have been developed that are based on aggregating events across time and multiple sources to identify network- or host-based behaviors that might be malicious. These behavior-based methods are not as brittle as signature-based techniques because they can detect new, as well as known, variations within a general class of attacks. Behavioral methods have been successful in finding previously undiscovered malware. However, most behavior-based detection tools are not real-time detectors. They require the development of robust classifiers that describe patterns of anomalous events representing potential misuse, ranging from low-level events such as the opening of a network connection to excessive Facebook use or watching World Cup soccer. Using these classifiers, behavior-based techniques typically find anomalies after the fact in batch-processed data. Anomalies are then ranked so that a human analyst can focus on the most significant problems. However, when an anomaly is determined to be part of a larger infection, these behavioral techniques produce important and unique signatures, which can then be used to stop infections in real time. More can be said about the current state of the art techniques in a closed session.

Current behavioral-based detection systems, however, are prone to high false positive rates. They require the supervision of skilled analysts to monitor and investigate alerts and to develop and adjust classifiers. The demand for skilled analysts far exceeds supply. Furthermore, difficult tasks can sometimes overwhelm even the best analysts. Depending on the time scale and complexity of the pattern of behavior associated with a particular type of malicious activity, behavioral techniques can also fail to detect an attack before an adversary has caused damage. Behavioral detection offers promise and will improve, but does not represent a panacea today.

An often overlooked component of cyber security is that anyone can obtain virtually any security product on the market. The fact that our adversaries can use their knowledge of common security tools to predict the barriers they might face during an attack suggests two requirements for network- and host-based intrusion detection systems: 1) signature-based products should provide an open interface by which we can develop and deploy proprietary signatures and scripts; 2) behavior-based tools that allow us to detect new attacks must be introduced to complement our signature-based methods. As behavioral-based detection systems improve, we anticipate a crossover where behavioral-based tools will become predominant and will be supplemented by signature-based methods.

4. Test and evaluation infrastructures at various classification levels (e.g. digital sandboxes): Experimentation plays a central role in science and engineering as a rigorous means of testing hypotheses and potential solutions. The cyber research and operational communities recognize the necessity of more realistic test and evaluation infrastructures, or test beds, to advance computer security research and conduct cyber planning, training, and exercises. Significant foundational work has been done through private-sector and government funded efforts, including the development of hardware and operating system emulation and virtualization tools, network traffic generators and

test bed management systems, and actual cyber test beds of varying size, realism, and classification levels. Examples include DoD Information Operations (IO) Range, and the National Cyber Range.

However, cyberspace is a highly complex, man-made environment of vast scale and heterogeneity and presents unique and daunting experimental challenges that we have not yet been able to adequately represent in test facilities. Our current capabilities fall short in fidelity and in scaling up to regional and Internet-sized networks. Additionally, while our adversaries use the Internet as their cyber test bed, it is not responsible for the United States to do the same because of possible, unintended side effects.

Sandia, in partnership with a number of government agencies and national laboratories, conducts significant research in cyber and cyber/physical test and evaluation technologies, including contributing roles in the IO Range, National Cyber Range, and DOE National Supervisory Control and Data Acquisition (SCADA) Test Bed. These activities build upon our long-standing investments and capabilities in high-performance computing and in modeling and simulation of physical and cyber systems. We and others have developed techniques and tools to conduct so-called live-virtual-constructive experiments that integrate real people and computer systems with simulated computer systems and modeled human behavior to evaluate consequences and mitigation strategies for realistic cyber scenarios like a cyber-attack on critical infrastructure.

Significant challenges remain, however, to realize the high-fidelity experiments required to support scientifically rigorous testing and evaluation of cyber solutions and scenarios. Cyber testing and evaluation can be broken down into four distinct experimental phases: design, configuration, execution, and result analysis. Research and development gaps remain in all four phases.

Cyber experiment design presents specific challenges stemming, in part, from the limited scientific foundation in cyber. In other disciplines, well-developed approaches like wind tunnel testing and scientific laws like those governing fluid dynamics can be brought to bear to design an effective experiment. By contrast, we struggle today to design good cyber experiments that are controlled and repeatable. The complexity from integrated circuits to Internet scale networks and the adversarial nature of cyberspace, make it difficult to design a complete, valid and meaningful experiment to study cyber phenomena of interest, such as the propagation of a botnet, or evaluate a prototype security technology. Additional work is needed to develop and promulgate a scientifically rigorous approach to designing cyber experiments and exercises.

There has been considerable progress in the last few years with tools and technologies for configuring and executing cyber experiments, but major gaps remain in these areas too. Although several test bed configuration tools now exist to specify and automatically configure elements like computer systems, and network topology, required for small experiments, large and complex experiments require time-consuming hand configuration and tuning of test bed elements. Configuration and execution of high fidelity, regional

and Internet-scale experiments still pose many research challenges. In some cases it is unclear what scale and fidelity are even needed to answer important questions.

Running realistically scaled experiments poses challenges of its own. Sandia recently demonstrated what we believe to be state-of-the-art scale by booting 4.5 million virtual computer nodes. These nodes were light-weight virtual machines, meaning they exhibit some, but not all, of the complex behavior of a typical desktop computer. However, at this scale one is getting close to representing the Internet resources of a small country. Current test beds also have overly simplistic human behavior modeling elements, and thus fail to adequately represent user frailties, like susceptibility to spear phishing - an e-mail spoofing fraud attempt that targets a specific organization, seeking unauthorized access to confidential data or the perverse creativity of adversaries.

The challenge of gathering and analyzing test results is also only partially solved. Fine-grained instrumentation is lacking from most existing test beds, as are tools for efficiently distilling and extracting pertinent results from the vast volumes of data that can be generated by large tests and exercises. Lastly, future test beds will need to be integrated in a much larger percentage of wireless components.

Advancing the state of the art in cyber test and evaluation will require major research and infrastructure investments. The government has already made large investments in this area through several standalone programs such as National Cyber Range. However, we see a need for a new strategy that coordinates future investments across the government in a way that maximizes technological advancements and ensures test bed access for academia, government, private-sector, and military users, while respecting agency- and program-specific test bed capability and classification requirements.

5. Other research priorities: We must devote additional attention to developing and implementing strategies for assuring the safety of the nation's most critical national security systems. These systems are particularly challenging to defend because of the full-spectrum attacks that a nation state or other highly capable threat actor is likely to employ.

The information technology supply chain is a particularly insidious risk to high-consequence national security systems, because of our widespread reliance on commercial-off-the-shelf (COTS) hardware and software technology that is increasingly produced in whole or in part by untrusted, non-US organizations. Unfortunately, the growing complexity of these systems also makes it economically infeasible to verify them thoroughly.

Insufficient attention has been given to technical approaches for mitigating supply chain risks. Counterfeiting and subversion of critical components in high-consequence DoD systems could have a devastating effect on our ability to project military power with confidence around the world. Better methodologies and technologies are needed for assessing and managing supply chain risks.

IT system trust must ultimately be rooted in hardware. Additional research is needed to enable scalable, cost-effective hardware integrity evaluation to verify that no malicious features have been added and that security features have not been weakened. We must be able to positively identify and track components throughout their complete lifecycle. We need to discover how to compose higher assurance systems from largely untrusted COTS components and a small set of simple trusted components.

To tip the balance in favor of defenders, we must create and deploy technologies and policies that decrease benefits and impose costs on attackers. Attackers are able to leverage the complexity of modern hardware and software systems to find and exploit a seemingly endless stream of vulnerabilities. These attacks scale globally to provide disproportionate benefit to attackers as a result of the relatively homogenous computing base that exists in most enterprise environments throughout the world. Although various secure design approaches, such as formal verification, offer promise, they do not currently scale to the size and complexity of COTS systems. In the near-term it is unlikely that COTS systems will be drastically simplified to facilitate formal methods-based, high-assurance development. Alternatively, approaches that introduce manageable and cost-effective diversity within hosts and across an enterprise could dramatically reduce the utility of many attacks and sharply raise development costs for attackers, forcing adversaries to have to discover and exploit multiple vulnerabilities simultaneously to mount a successful attack.

6. **Workforce issues:** Confronting the challenges I have outlined today requires a highly skilled and motivated research community. It is well documented that the demand for cyber expertise greatly exceeds the supply.[1,2] Over the past three years Sandia has been able to attract and hire top United States citizen undergraduate talent by paying for their master's degree at the school of their choice and supporting them with 75% of their salary while they attend school full time. Upon returning to Sandia, they owe us two years without penalty. This has been a very successful recruiting program but retention results won't be available for a few more years. Doctoral and experienced cyber hires are more difficult, even with market-based salary offers, because of intense competition for their knowledge and skills. However, we have been successful in attracting a few high-quality Ph.D. researchers through a new competitive early-career research program that provides selected Ph.D. hires with two years of internal funding for independent research.

Retention is a growing concern. Although the importance of the national security mission and job stability remain highly attractive features to our employees, new hires today receive benefits similar to those found in U.S. industry. Over time, therefore, we may see the retention rate for computer science professional's approach that of industry, which retains such staff for approximately five years. This could become a significant issue because it takes three to five years of mentoring for a recent graduate to become highly skilled in supporting cyber research for the U.S. government.

[1] *http://www.cioinsight.com/c/a/Trends/Damn-the-Economy-IT-Employment-Rises-to-New-Heights/*

[2] *Langevin Assesses State of Cyber Workforce, http://langevin.house.gov/news/press-releases/2011/10/langevin-assesses-state-of-cyber-workforce.shtml*

Historically, the laboratories are asked to solve the "impossible" problems. Congress should consider the implications of not having the best and brightest U.S. cleared and experienced staff available to tackle the nation's most challenging security needs. Presently, many of Sandia's cyber staff are being solicited by private companies offering more than 50% increases in salary and better benefits. Historically, we have lost less than a percent of our cyber workforce to outside employment; however, we are currently on a path to lose 10% this fiscal year.

Outside of the labs' recruitment and retention challenges, there are additional areas that deserve attention. Academic programs for computer security specializations need improvement. Curricula vary from one university to another and few programs produce graduates who have both the required deep knowledge of computer hardware and systems combined with practical security understanding and skills. The Scholarship For Service (SFS) program has helped produce more qualified graduates, but in my opinion could be enhanced to attract the nation's best students who are in turn intentionally cultivated for government service through improved curricula and hands-on training programs. Government labs and agencies participate today by providing SFS students with internships and hiring SFS graduates, but we could also partner with SFS-funded universities to help develop appropriate curricula, training toolkits, and exercises.

Beyond SFS, the labs can serve a broader role as a training ground for the nation's next generation of security researchers and operational defenders. For the past 10 years Sandia has run an innovative hands-on computer security internship program for undergraduate and graduate students called the Center for Cyber Defenders (CCD). Drawing summer projects from our customer-funded security R&D programs provides students with an opportunity to work on real security problems and experience the satisfaction of contributing directly to national security. For the first time this year, thanks to Department of Homeland Security (DHS) S&T support, we will be piloting a secure systems research challenge for CCD students that we hope can be extended to include other labs. In general, we believe student competitions are an important and still underutilized mechanism to attract, engage, and accelerate the development of cyber professionals.

Professional education and training is another challenge. Knowledge in cyber disciplines constantly evolves, often in obscure corners of the Internet. Continuous learning and skills refreshing are required to maintain a world-class R&D and operational cyber workforce. We and others have done some preliminary work on competency-based training and other professional development activities such as rotational assignments between research and mission-focused roles, but this area requires additional attention, especially in light of the magnitude of the government's cyber workforce needs and the retention issues mentioned previously.

7. **Current coordination across the community:** From a laboratory R&D perspective, coordination is good. For example, DoD T&E reaches out to the labs that have specific skills and the labs coordinate well with each other in assessing and improving DoD IT

systems. Coordination is similarly close with other government agencies including people working together at each other's sites and through quarterly reviews.

From an operational perspective, coordination within the federal government is improving. US-CERT has created capable collaboration facilities within their secure web site. In our opinion there is still too much focus on security compliance. Compliance-based security is not effective. When coupled with excessive oversight, a compliance focus results in brittle and unresponsive security systems. Today, victims are often punished for the actions of adversaries.

SUMMARY AND CONCLUSIONS: To tip the balance in favor of defenders, approaches and technologies must be developed and deployed that decrease benefits and impose costs (or risk) to attackers. Attackers are able to leverage the complexity of modern hardware and software systems at the component level to find and exploit a seemingly endless stream of vulnerabilities. These attacks scale globally to provide disproportionate benefit to attackers as a result of the relatively homogenous computing base that exists in most enterprise environments throughout the world. However, the cost equation to the adversary can be changed. Cyber defensive technology has been shown to accelerate when long-term stable funding is in place, technical collaboration among research organizations involves "prisoner exchanges," test facilities are prepositioned and analysis/operators are an integral part of the team. As one example, behavioral-based detection systems are having significant success and as they improve, eventually we anticipate a crossover where behavioral-based tools will become predominant and supplemented by signature-based methods.

Two areas within the scope of this Committee's questions need to be addressed: 1) the test environments available to the research community; and 2) the retention of the government's cyber research community, which includes the national laboratories. To continue the acceleration of government-developed and-owned cyber defense technologies, testing and emulation environments of various combinations of scale, fidelity, and heterogeneous representations of regional and Internet-sized networks are needed to address multiple national security missions. With their deep reservoir of technical talent and science and technology capabilities, the DOE national laboratory complex has helped address some of the government's most challenging national security problems, including cyber. However, unlike the Cold War where the government used work environment, benefits and mission to attract and retain top scientists to government agencies and national labs, only a small fraction of those retention tools exist for the cyber war and the implications should be of great concern.

B I O G R A P H Y

Dr. James Peery

Director of Information Systems
Sandia National Laboratories

Dr. James Peery is the Director of the Information Systems Analysis Center, at Sandia National Laboratories (SNL) in Albuquerque, New Mexico. In this role, he is responsible for development and application of new information technologies that enable information superiority for national security and critical infrastructure protection customers. In addition, James is the Program Director for SNL's Information Operations Program.

From 2007 to March 2010, James was the Director of the Computation, Computers, Information and Mathematics (CCIM) Center. CCIM is the foundation of SNL's research and development activities in high performance computing. CCIM contains the Computer Science Research Institute (CSRI), the joint Institute for Advanced Architectures and Algorithms (IAA) with ORNL and the Alliance for Computing at Extreme Scales (ACES) with LANL. During this period, James was the Program Director of the NNSA's Advanced Simulation and Computing Program at SNL.

Prior to returning to Sandia, James worked at Los Alamos National Laboratory (LANL) from 2002 to 2007 in the positions of Hydrodynamic Experiments Division Leader, Principal Deputy Associate Director of the LANL's Nuclear Weapons program and Program Director of the NNSA's Advanced Simulation and Computing Program. Before joining LANL, James worked at Sandia National Laboratories where he led the Computational Solid Mechanics and Structural Dynamics Department and Computation Physics Department. During his career, James has been responsible for the development of state-of-the-art, massively parallel computational tools in the fields of high energy density physics, shock physics, transient dynamics, quasistatics, nonlinear implicit dynamics, and structural dynamics. James' major research areas are in Arbitrary Lagrangian Eulerian (ALE) algorithms and parallel algorithms where he has published greater than 50 papers. As part of the SALINAS team, James was awarded the 2002 Gordon Bell Award and NNSA Award for Excellence. James earned his Ph.D. degree in nuclear engineering from Texas A&M University and joined Sandia National Laboratories as a Member of the Technical Staff in 1990.

U.S. DEPARTMENT OF ENERGY NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2011-XXXXP

Sandia National Laboratories

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2011-XXXXP